



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/726,948	12/02/2003	Stephen C. Hahn	SUNMP453	5353
32291	7590	05/12/2008	EXAMINER	
MARTINE PENILLA & GENCARELLA, LLP			WAI, ERIC CHARLES	
710 LAKEWAY DRIVE				
SUITE 200			ART UNIT	PAPER NUMBER
SUNNYVALE, CA 94085			2195	
			MAIL DATE	DELIVERY MODE
			05/12/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/726,948	HAHN ET AL.	
	Examiner	Art Unit	
	ERIC C. WAI	2195	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 18 January 2008.

2a) This action is **FINAL**. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-53 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) Claim(s) _____ is/are allowed.

6) Claim(s) 1-53 is/are rejected.

7) Claim(s) _____ is/are objected to.

8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a) All b) Some * c) None of:

1. Certified copies of the priority documents have been received.
2. Certified copies of the priority documents have been received in Application No. _____.
3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) <input type="checkbox"/> Notice of References Cited (PTO-892)	4) <input type="checkbox"/> Interview Summary (PTO-413)
2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)	Paper No(s)/Mail Date. _____ .
3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)	5) <input type="checkbox"/> Notice of Informal Patent Application
Paper No(s)/Mail Date _____ .	6) <input type="checkbox"/> Other: _____ .

DETAILED ACTION

1. Claims 1-53 are presented for examination.

Claim Rejections - 35 USC § 103

2. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

3. Claims 1-35, and 38-53 are rejected under 35 U.S.C. 103(a) as being unpatentable over Applicant's Admitted Prior Art (AAPA) in view of Uri ("Workload management-many questions", IT Resource Forums, 7-10-2002, XP002290900, pp, 1-4, <<http://forums1.itrc.hp.com/service/forums/questionanswer.do?admit+716493758+1091180233157+28353475&threadID+25550>>).

4. Uri was disclosed in IDS dated 06/20/2005.

5. Regarding claim 1, AAPA teaches a computer implemented method of managing computer resources comprising:

accessing an amount of computer resources allocated to a workload of a computer system, the workload exists within a user space and includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space (pg 3 lines 14-16, wherein resources are accessed; col 3

line 22- col 4 line 5, wherein workloads exist in a user space and include child processes);

monitoring computer resource usage of said workload and only for resources associated with said workload(pg 3 lines 14-16, wherein the request is determined to request more physical memory than is available); and

determining a range of computer resources to make available for other use by other workloads provided said computer resource usage of said workload exceeds said amount of computer resources allocated to said workload, the computer resources being made available to other workloads by paging out least recently used physical memory assigned to said workload (pg 3 lines 16-20 and pg 5 line 11-12, wherein in response to the overcommit, the system will free up resources for the requestor; pg 4 lines 20-22, wherein the memory to be paged out is least recently used).

6. AAPA does not explicitly teach wherein said monitoring is performed from within said workload and said monitoring and determining occur within the user space. In fact, AAPA teaches the steps of monitoring and determining for all resources (i.e. all user spaces) by an operating system.

7. Uri teaches an improvement on existing techniques that is supported by HP-UX 11i. Uri teaches using Memory Resource groups wherein a block of memory is allocated to each workload and a separate memory management subsystem will be created for each workload, i.e. part of that workload (pg 3 paragraph 2). Uri further teaches that if an application is a workload attempts to allocate more memory than it is entitled to,

paging will occur despite there being addition memory being available to other workloads (pg 3 paragraph 2).

8. It would have been obvious to one of ordinary skill in the art at the time of the invention that the accessing is performed within the user space. One would be motivated by the desire to make these modifications to AAPA because doing so allows for more fine grained control and allows for resources to be allocated based on business priorities as indicated by Uri (pg 3 paragraph 3).

9. Regarding claim 2, AAPA teaches that the computer resources comprise physical memory (pg 3 lines 14-15).

10. Regarding claim 3, AAPA teaches that said determining comprises determining if a page of physical memory utilized by said workload has been accessed by said workload within a predetermined period of time (pg 4 lines 20-22).

11. Regarding claim 4, AAPA and Uri do not explicitly teach determining if a page of physical memory utilized by said workload has been accessed by said workload within a predetermined period of time comprises determining if said page has been accessed by said workload since a previous determination of whether said page had been accessed by said workload.

12. It would have been obvious to one of ordinary skill in the art at the time of the invention, to repeatedly make the determination whether a page in the physical memory

has been recently accessed. One would be motivated by the desire to continuously monitor the usage of each page within the physical memory.

13. Regarding claim 5, AAPA teaches that computer resources comprises virtual memory (pg 3 lines 23-24).

14. Regarding claims 6-8, AAPA and Uri do not teach that the computer resources comprises central processing unit time, input/output space, or network bandwidth.

15. However, it would have been obvious to one of ordinary skill in the art at the time of the invention to include computer resources comprising central processing unit time, input/output space, or network bandwidth. It is well known in the art that computers are constrained by such resources.

16. Regarding claim 9, AAPA teaches a computer implemented method comprising:
monitoring usage of a computing resource utilized by a workload, the workload includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space (pg 3 lines 14-16, wherein the request is determined to request more physical memory than is available; col 3 line 22-col 4 line 5, wherein workloads exist in a user space and include child processes);
responsive to exceeding a limit on utilization of said computing resource, decreasing usage of said computing resource by said workload (pg 5 lines 1-9, wherein a given workload is victimized when it exceeds its allocation).

17. AAPA does not explicitly teach wherein said monitoring occurs within a user space and the process monitors the user space only. In fact, AAPA teaches the step of monitoring for all resources (i.e. all user spaces) by an operating system.

18. Uri teaches an improvement on existing techniques that is supported by HP-UX 11i. Uri teaches using Memory Resource groups wherein a block of memory is allocated to each workload and a separate memory management subsystem will be created for each workload, i.e. part of that workload (pg 3 paragraph 2). Uri further teaches that if an application is a workload attempts to allocate more memory than it is entitled to, paging will occur despite there being addition memory being available to other workloads (pg 3 paragraph 2).

19. It would have been obvious to one of ordinary skill in the art at the time of the invention that the accessing is performed within the user space. One would be motivated by the desire to make these modifications to AAPA because doing so allows for more fine grained control and allows for resources to be allocated based on business priorities as indicated by Uri (pg 3 paragraph 3).

20. Regarding claim 10, AAPA teaches that said computing resource comprises physical memory (pg 3 lines 14-15).

21. Regarding claim 11, AAPA teaches that decreasing usage of said computing resource comprises paging a portion of said physical memory assigned to said workload out of said physical memory (pg 4 lines 7-13).

22. Regarding claim 12, AAPA teaches that said portion of said physical memory comprises a least recently used portion of said physical memory assigned to said workload (pg 4 lines 20-22).

23. Regarding claim 13, AAPA teaches that said decreasing usage does not halt operation of said workload (pg 5 lines 8-9, wherein the workload will continue operation).

24. Regarding claim 14, AAPA teaches that said decreasing usage is initiated by a process of said workload (pg 5 lines 18-20).

25. Regarding claim 15, AAPA and Uri do not teach that said process that performs said monitoring is not an operating system kernel process.

26. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

27. Regarding claim 16, AAPA teaches a computer implemented method for memory management of a workload comprising:

accessing a list of memory pages assigned to said workload, the workload includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space (pg 5 lines 11-12 and col 3 line 22-col 4 line 5, wherein workloads exist in a user space and include child processes);

responsive to a request from a first process of said workload for memory which exceeds a predetermined memory limit for said workload, selecting a plurality of memory pages from said list of memory pages, wherein the plurality of memory pages includes least recently used memory pages assigned to the workload (pg 5 lines 11-12 and pg 4 lines 20-22, wherein the memory to be paged out is least recently used); and

initiating a second process within the user space to page out said plurality of memory pages (pg 5 lines 11-13).

28. AAPA does not explicitly teach that the accessing is performed in a user space. However, Uri teaches an improvement on existing techniques that is supported by HP-UX 11i. Uri teaches using Memory Resource groups wherein a block of memory is allocated to each workload and a separate memory management subsystem will be created for each workload, i.e. part of that workload (pg 3 paragraph 2). Uri further teaches that if an application is a workload attempts to allocate more memory than it is entitled to, paging will occur despite there being addition memory being available to other workloads (pg 3 paragraph 2).

29. It would have been obvious to one of ordinary skill in the art at the time of the invention that the accessing is performed within the user space. One would be motivated by the desire to make these modifications to AAPA because doing so allows for more

fine grained control and allows for resources to be allocated based on business priorities as indicated by Uri (pg 3 paragraph 3).

30. Regarding claim 17, AAPA teaches that accessing, selecting and initiating are performed by said second process within said workload (pg 5 lines 11-20, wherein the steps are performed by another process).

31. Regarding claim 18-19, AAPA and Uri do not teach that said second process is not an operating system kernel process or that said second process is loaded into a user space.

32. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

33. Regarding claims 20-21, AAPA teaches that said plurality of memory pages comprises memory pages that are least recently used by said workload (pg 4 lines 20-23).

34. Regarding claim 22, AAPA and Uri do not explicitly teach that said page out of said plurality of least recently used memory pages reduces a number of memory pages assigned to said workload to below said memory limit.

35. However, it would have been obvious to one of ordinary skill in the art at the time of the invention, that if the page out was performed in response to the physical requirements being exceeded, that the course of action would remedy the excessive use of resources.

36. Regarding claim 23, AAPA and Uri do not teach that said plurality of least recently used memory pages comprises the minimum number of memory pages to reduce said number of memory pages assigned to said workload below said memory limit.

37. However, it would have been obvious to one of ordinary skill in the art at the time of the invention, that if the page out was performed in response to the physical requirements being exceeded, that the course of action would remedy the excessive use of resources.

38. Regarding claim 24, AAPA teaches wherein at least a portion of said workload continues to operate subsequent to said initiating (pg 5 lines 8-9, wherein the workload will continue operation).

39. Regarding claim 25, AAPA and Uri do not teach that said initiating is not performed by an operating system kernel process.

40. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

41. Regarding claims 26-31, they are the computer-usable medium claims of claims 1-4 above. Therefore, they are rejected for the same reasons as claims 1-4 above.

42. Regarding claim 32, AAPA teaches a computer implemented method comprising: scanning pages for a first workload of a computer to determine if each of said pages was accessed since a last scan (pg 4 lines 20-23, wherein the selecting checks all memory pages from throughout the entire computer system); repeating said scanning and said setting for a second workload (pg 4 lines 20-23, wherein the selecting checks all memory pages from throughout the entire computer system),

wherein the first workload and the second workload exist within the user space and each includes a plurality of running processes, the plurality of running processes are a subset of all processes the are running in the user space (col 3 line 22- col 4 line 5, wherein workloads exist in a user space and include child processes).

43. AAPA does not teach setting bits indicative of the result of said scanning within a scoreboard related to said first workload. However, it would have been obvious to one of ordinary skill in the art at the time of the invention to utilize some scoreboard or table to quickly assess which pages were the least recently used.

44. AAPA also does not teach that the scanning and setting occur within a user space of said computer.

45. Uri teaches an improvement on existing techniques that is supported by HP-UX 11i. Uri teaches using Memory Resource groups wherein a block of memory is allocated to each workload and a separate memory management subsystem will be created for each workload, i.e. part of that workload (pg 3 paragraph 2).

46. It would have been obvious to one of ordinary skill in the art at the time of the invention that the scanning and setting occur within the user space. One would be motivated by the desire to make these modifications to AAPA because doing so allows for more fine grained control and allows for resources to be allocated based on business priorities as indicated by Uri (pg 3 paragraph 3).

47. Regarding claim 33, AAPA and Uri do not teach scanning comprises checking said pages according to an order inherent to a list of said pages.

48. However, it is well known in the art to scan a list in order.

49. Regarding claim 34, AAPA teaches paging out a plurality of pages utilized by said first workload responsive to said determining (pg 4 lines 20-22).

50. Regarding claim 35, AAPA teaches determining if the number of pages utilized by said workload exceeds a predetermined limit (pg 4 lines 7-10).

51. Regarding claim 38, AAPA teaches a computer implemented method of managing computer resources over a plurality of workloads, said method comprising:

for each workload of said plurality of workloads, monitoring respective workload resource usage against a respective allotment of each workload (pg 5 lines 11-12);

determining a range of computer resources to page out for each workload whose resource usage exceeds its respective allotment (pg 5 lines 11-12); and

initiating a paging out operation of said range of computer resources and wherein said monitoring (pg 5 lines 11-13),

wherein each of the plurality of workloads exists within a user space and includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space (pg 5 lines 11-12 and col 3 line 22- col 4 line 5, wherein workloads exist in a user space and include child processes),

52. AAPA does not explicitly teach wherein said determining and initiating occur within a user space. In fact, AAPA teaches the step of selecting for all resources (i.e. all user spaces).

53. Uri teaches an improvement on existing techniques that is supported by HP-UX 11i. Uri teaches using Memory Resource groups wherein a block of memory is allocated to each workload and a separate memory management subsystem will be created for

each workload, i.e. part of that workload (pg 3 paragraph 2). Uri further teaches that if an application is a workload attempts to allocate more memory than it is entitled to, paging will occur despite there being addition memory being available to other workloads (pg 3 paragraph 2).

54. It would have been obvious to one of ordinary skill in the art at the time of the invention that the determining and initiating occur within a user space. One would be motivated by the desire to make these modifications to AAPA because doing so allows for more fine grained control and allows for resources to be allocated based on business priorities as indicated by Uri (pg 3 paragraph 3).

55. Regarding claim 39, AAPA teaches that determining comprises determining least recently used pages for each workload whose resource usage exceeds its respective allotment (pg 5 lines 11-13).

56. Regarding claim 40, AAPA teaches that the process is situated within a workload of said plurality of workloads (pg 3 lines 22-24).

57. Regarding claim 41, AAPA teaches paging out said range of computer resources and wherein each workload whose resource usage exceeds its respective allotment remains partially operable during said paging out of its respective range of computer resources (pg 5 lines 1-9).

58. Regarding claim 42, AAPA and Uri substantially teach the claim according to the reasons given in claim 1 above. However, AAPA and Uri fail to teach a bus for functionally coupling elements of said computer system; physical memory coupled to said bus for storing processor instructions and data; and a processor coupled to said bus.

59. However, official notice is made that it is well known in the art that a computer system consists of a processor, memory, and buses interconnecting elements.

60. Regarding claims 43-47, they are the computer system claims of claims 2, and 5-8 above. Therefore, they are rejected for the same reasons as claims 2, and 5-8.

61. Regarding claim 48, AAPA and Uri substantially teach the claim according to the reasons given in claim 1 above. However, AAPA and Uri fail to teach that the steps of accessing, monitoring, and determining are performed by a first computer on workloads belonging to a second computer.

62. AAPA teaches that such steps are performed by the operating system but that such operations are costly in terms of computing load (pg 5 lines 12-15). Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to offload the monitoring steps to another computer. The skilled artisan would realize that doing so would free resources that would then be available to process the loads. One would be motivated by the desire to reduce the computing load on a single system.

63. Regarding claims 49-53, they are the computer system claims of claims 2, and 5-8 above. Therefore, they are rejected for the same reasons as claims 2, and 5-8.

64. Claims 36-37 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tanenbaum (Modern Operating Systems, 2nd edition, 2001, Prentice Hall Intl, New Jersey) in view of Applicant's Admitted Prior Art (AAPA).

65. Tanenbaum was disclosed in IDS dated 06/20/2005.

66. Regarding claim 36, Tanenbaum teaches a computer implemented method comprising:

accessing memory usage for a workload and examining page usage for each process of said workload (pg 234, “4.6.1 Local versus Global Allocation Policies”, paragraphs 2-3, wherein the algorithm tries to find the least recently used page for all the processes);

aggregating usage of said each process to determine an aggregate usage for said workload (wherein it is inherent that a workload comprises many processes);

determining least recently used pages based on accessed bits associated with said workload (pg 235, Fig 4-28, wherein it is inherent that bits are set to indicate least recently used pages).

67. Tanenbaum does not teach performing the determining step based on whether the aggregate usage exceeds said memory utilization limit for said workload and

supplying a range of least recently used pages in a system call to an operating system kernel for evicting said range of least recently used pages to reduce resource usage by said workload; and retaining at least partial operation of said workload during said page evicting, wherein the workload exists within a user space and includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space.

68. AAPA teaches paging out memory in response to workloads exceeding their respective allotments whereby an operating system kernel process will evict selected pages (pg 5 lines 11-20) and that the workload exists within a user space and includes a plurality of running processes, the plurality of running processes are a subset of all processes that are running in the user space (pg 5 lines 11-12 and col 3 line 22- col 4 line 5, wherein workloads exist in a user space and include child processes)

69. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the determination in response to an aggregate usage exceeding the memory utilization limit. One would be motivated by the desire to ensure that overall memory usage does not exceed the amount that is physically available.

70. AAPA teaches that said determining and said supplying occur in a plurality of user space processes (pg 3 line 24 to pg 4 line 1, wherein multiple user spaces can exist in a system).

Response to Arguments

71. Applicant's arguments with respect to claims 1, 9, 16, 26, 32, 38, 42, and 48 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

72. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

73. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Eric C. Wai whose telephone number is 571-270-1012. The examiner can normally be reached on Mon-Thurs, 9am-5pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Meng - Ai An can be reached on 571-272-3756. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Eric C Wai/
Examiner, Art Unit 2195

/Meng-Ai An/
Supervisory Patent Examiner, Art Unit 2195